IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicant(s): Cain Application No.: 09/457209 Filed: 12/08/1999 Title: System, Device, and Method for Sending Keep-Alive Messages in a Communication Network Attorney Docket No.: 120-025	Group Art Unit: 2431 Examiner: Zia

Mail Stop Appeal Brief-Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

APPEAL BRIEF

Sir:

Please enter this Appeal Brief in response to the Notice dated July 6, 2011.

I. Real Party in Interest

The real party in interest is Nortel Networks Limited.

II. Related Appeals and Interferences

Appellants are not aware of any related appeals or interferences.

III. Status of the Claims

Claims 1, 2, 4-9, 11-16, 18-22 and 24-26 are pending in this application. All of the pending claims are rejected. No claims have been allowed. The rejections of claims 1, 2, 4-9, 11-16, 18-22 and 24-26 are the subject of this appeal.

IV. Status of Amendments

All submitted amendments have been entered and considered.

V. Summary of Claimed Subject Matter

The presently claimed invention relates generally to communication systems, and more particularly to sending keep-alive messages in a communication system. Computers and computer peripherals are often internetworked over a communication network. The communication network includes a number of network nodes that interoperate to route protocol messages within the communication network. These network nodes typically run various routing protocols in order to determine forwarding paths for routing protocol

messages within the communication network. When a network node fails, the other network nodes need to route the protocol messages around the failed network node. The network nodes typically rely on "keepalive" messages to determine whether a particular network node is operational. Each node periodically sends keep-alive messages to its neighbors. A network node may consider a particular neighbor to be operational as long as the neighbor is sending keep-alive messages. Therefore, each network node receives keep-alive messages from its neighbors. The processing of keep-alive messages can be computationally intensive, especially if the network node has many neighbors.

In accordance with one aspect of the invention, the frequency for sending keepalive messages to a neighbor is determined based upon a reliability factor for communicating with the neighbor. A node determines a reliability factor for communicating with a neighbor and sets the frequency for sending keep-alive messages to the neighbor based upon the reliability factor. The reliability factor is determined based upon the reliability of the neighbor as well as the reliability of the communication link to the neighbor. The frequency for sending keep-alive messages to the neighbor is relatively high if the reliability factor is low. The frequency for sending keep-alive messages to the neighbor is relatively low if the reliability factor is high. The frequency for sending keepalive messages to the neighbor is dynamically adjusted based upon an updated reliability factor.

The limitations recited in the independent claims are supported by the specification and drawing as indicated in bold below.

1. (previously presented) A method for sending keep-alive messages by a node to a neighbor in a communication network, the method comprising:

measuring a reliability of a communication link to the neighbor; and

The reliability factor is preferably based upon the reliability of the neighbor as well as the reliability of the communication link to the neighbor, both of which can be determined to some degree and can vary over time. Page 3:14-16

periodically calculating a reliability factor for communicating with a neighbor based upon the reliability of the communication link to the neighbor;

Each node determines a reliability factor for communicating with each of its neighbors. The node continually updates the reliability factor and adjusts the frequency of keep-alive messages accordingly. Page 3:12-18

varying a frequency for sending keep-alive messages to the neighbor based upon the reliability factor; and

An exemplary embodiment of the present invention sets the frequency for sending keep-alive messages to a particular neighbor based upon a reliability factor for communicating with the neighbor. The keep-alive messages are sent at a relatively low frequency if the reliability factor for communicating with the neighbor is high. This is because there is a relatively high likelihood that each keep-alive message will be received and processed by the neighbor, so fewer keep-alive messages are needed to keep the communication link to the neighbor active. The keep-alive messages are sent at a relatively high frequency if the reliability factor for communicating with the neighbor is low. This is because there is a relatively low likelihood that each keep-alive message will be received and processed by the neighbor, so more keep-alive messages are needed to keep the communication link to the neighbor active. The reliability factor is updated regularly, and the frequency for sending keep-alive message to the neighbor is dynamically adjusted accordingly. Page 2:31 through page 3:11

sending keep-alive messages by the node to the neighbor in accordance with those steps. **Id.**

8. (previously presented) A device for sending keep-alive message to a neighbor in a communication network, the device comprising:

a processor, a memory and a transmitter;

In an exemplary embodiment of the present invention, predominantly all of the logic for sending keep-alive messages described herein is implemented as a set of computer program instructions that are stored in a computer readable medium and executed by an embedded microprocessor system within a network node. Page 5:29 through page 6:2.

the processor executing a computer program stored in the memory, the computer program including:

reliability calculation logic operably coupled to determine a reliability for a communication link to the neighbor and periodically calculate a reliability factor for communicating with the neighbor based upon the reliability for the communication link to the neighbor; and

The reliability factor is preferably based upon the reliability of the neighbor as well as the reliability of the communication link to the neighbor, both of which can be determined to some degree and can vary over time. Each node determines a reliability factor for communicating with each of its neighbors. The node continually updates the reliability factor and adjusts the frequency of keep-alive messages accordingly. Page 3:12-18

frequency variation logic responsive to the reliability calculation logic and operably coupled to calculate a frequency for sending keep-alive messages to the neighbor via the transmitter based upon the reliability factor.

An exemplary embodiment of the present invention sets the frequency for sending keep-alive messages to a particular neighbor based upon a reliability factor for communicating with the neighbor. The keep-alive messages are sent at a relatively low frequency if the reliability factor for communicating with the neighbor is high. This is because there is a relatively high likelihood that each keep-alive message will be received and processed by the neighbor, so fewer keep-alive messages are needed to keep the communication link to the neighbor active. The keep-alive messages are sent at a relatively high frequency if the reliability factor for communicating with the neighbor is low. This is because there is a

relatively low likelihood that each keep-alive message will be received and processed by the neighbor, so more keep-alive messages are needed to keep the communication link to the neighbor active. The reliability factor is updated regularly, and the frequency for sending keep-alive message to the neighbor is dynamically adjusted accordingly. Page 2:31 through page 3:11

15. (previously presented) A program product recorded on a computer readable medium for sending keep-alive messages to a neighbor in a communication network, the program product comprising:

reliability calculation logic operably coupled to measure a reliability for a communication link to the neighbor and to periodically calculate a reliability factor for communicating with the neighbor based upon the reliability for the communication link to the neighbor; and

The reliability factor is preferably based upon the reliability of the neighbor as well as the reliability of the communication link to the neighbor, both of which can be determined to some degree and can vary over time. Each node determines a reliability factor for communicating with each of its neighbors. The node continually updates the reliability factor and adjusts the frequency of keep-alive messages accordingly. Page 3:12-18

frequency variation logic responsive to the reliability calculation logic and operably coupled to determine a frequency for sending keep-alive messages to the neighbor based upon the reliability factor.

An exemplary embodiment of the present invention sets the frequency for sending keep-alive messages to a particular neighbor based upon a reliability factor for communicating with the neighbor. The keep-alive messages are sent at a relatively low frequency if the reliability factor for communicating with the neighbor is high. This is because there is a relatively high likelihood that each keep-alive message will be received and processed by the neighbor, so fewer keep-alive messages are needed to keep the communication link to the neighbor active. The keep-alive messages are sent at a relatively high frequency if the reliability factor for communicating with the neighbor is low. This is because there is a relatively low likelihood that each keep-alive message will be received and processed by the neighbor, so more keep-alive messages are needed to keep the communication link to the neighbor active. The reliability factor is updated regularly, and the frequency for sending

keep-alive message to the neighbor is dynamically adjusted accordingly. Page 2:31 through page 3:11

22. (previously presented) Apparatus comprising:

a plurality of interconnected devices including a node and a neighbor in communication over a link,

wherein the node is operably coupled to send keep-alive messages to the neighbor, and wherein the node is operably coupled to vary the frequency for sending keep-alive messages to the neighbor based upon a periodically computed reliability factor for communicating with the neighbor over the communication link, wherein the node is operably coupled to calculate the reliability factor based upon a reliability for the neighbor and a measured reliability for the communication link.

An exemplary embodiment of the present invention sets the frequency for sending keep-alive messages to a particular neighbor based upon a reliability factor for communicating with the neighbor. The keep-alive messages are sent at a relatively low frequency if the reliability factor for communicating with the neighbor is high. This is because there is a relatively high likelihood that each keep-alive message will be received and processed by the neighbor, so fewer keep-alive messages are needed to keep the communication link to the neighbor active. The keep-alive messages are sent at a relatively high frequency if the reliability factor for communicating with the neighbor is low. This is because there is a relatively low likelihood that each keep-alive message will be received and processed by the neighbor, so more keep-alive messages are needed to keep the communication link to the neighbor active. The reliability factor is updated regularly, and the frequency for sending keep-alive message to the neighbor is dynamically adjusted accordingly. Page 2:31 through page 3:11 The reliability factor is preferably based upon the reliability of the neighbor as well as the reliability of the communication link to the neighbor, both of which can be determined to some degree and can vary over time. Each node determines a reliability factor for communicating with each of its neighbors. The node continually updates the reliability factor and adjusts the frequency of keep-alive messages accordingly. Page 3:12-18

VI. Grounds of Rejection to be Reviewed on Appeal

A. Claims 1, 2, 4-9, 11-14, 22, and 24-26 are rejected on the grounds of nonstatutory double patenting over claims 1-26 of U.S. 7,035,214 since the claims, if allowed, would allegedly improperly extend the "right to exclude" already granted in the patent

VII. Argument

A. The examiner has not applied the correct standard, and the pending claims are patentably distinct from the claims of the '214 patent.

The '214 patent describes transmitting data in a data communications network using a transmission control protocol in a manner which reduces acknowledgment control traffic, and improved error recovery and congestion control. Claim 1 of the '214 patent is exemplary. It recites determining, at the transmitter, if an acknowledgment to the keep-alive request is not received before expiry of the re-transmission time-out timer, whereupon the transmitter backs off for a predetermined period; detecting a missing data packet at the receiver; sending a negative acknowledgment from the receiver to the transmitter for the missing data packet, the receiver being unresponsive to any packets from the

transmitter unless the receiver detects the missing data packet; and decreasing, at the transmitter, the length of the congestion window in response to receipt of the negative acknowledgment. The specification describes the same operating steps and states that the congestion window determines the transmission rate. Claims 10, 14, 20, 24, and 25 recite corresponding limitations. It is therefore undisputable that the claims of '214 patent recite *decreasing transmission rate in response to negative acknowledgements*.

Claim 1 of this application recites measuring reliability of a communication link to the neighbor; periodically calculating a reliability factor for communicating with the neighbor based upon the reliability of the communication link to the neighbor; varying the frequency for sending keep-alive messages to the neighbor based upon the reliability factor; and sending keep-alive messages by the node to the neighbor in accordance with those steps. In other words, the frequency of sending keep-alive messages to the neighbor is set as a function of reliability of communication with the neighbor. As described in the Abstract, page 3 and elsewhere, the frequency of sending keep-alive messages to a reliable neighbor may be lower than the frequency of sending keep-alive messages to an unreliable neighbor. Consequently, resources used to detect a failure would be related to likelihood of failure. Furthermore, the failure of a relatively frequently failing node would be detected more quickly. Claims 8, 15, and 22 recite corresponding limitations. Claims 2, 4-7, 9, 11-14, and 24-26 further define their respective base claims.

Comparing the pending claims in this application with the issued claims in the '214 patent it is clear that the pending claims are patentably distinct from the issued claims. The '214 patent claims a technique based on *negative acknowledgements from the receiver* as opposed to *keep alive messages from the transmitter*. Further, the claims of the '214 patent recite a *congestion control* technique based on *changing transmission rate of (non-management) data packets*, whereas the pending claims recite *resource management* based on *changing the frequency of keep alive messages (management packets)* based on reliability of the receiver, resulting in *use of failure detection resources commensurate with likelihood of failure*. Both sets of claims describe communication between network nodes. However, the recited claim elements and results are fundamentally different. Appellant therefore asserts that the pending claims are patentably distinct from the claims of the '214 patent.

In the Final Office Action dated December 27, 2010 the examiner asserts that the pending claims in this application and the claims of the '214 patent "both … claim subject matter that relates to … transmitting data in a data communications network, using a transmission control protocol, to provide reduced and adjusted acknowledgment control traffic, and both … [reduce] the acknowledgment traffic generated by TCP and control error recovery and congestion that does not require acknowledgements." Based on those assertions the examiner concludes that "the cited prior art does teach or suggest the subject matter broadly recited in [the pending claims]." Appellant asserts that the examiner has not applied the correct standard, and has therefore erred as a matter

of law. The test for double-patenting is not whether the claims of the prior art teach or suggest the subject matter broadly recited in the pending claims, but rather whether an examined application claim is not patentably distinct from the reference claim. See, e.g., *In re Berg*, 140 F.3d 1428, 46 USPQ2d 1226 (Fed. Cir. 1998). The analysis employed in an obviousness-type double patenting rejection parallels the guidelines for analysis of a 35 U.S.C. 103 obviousness determination. *In re Braat*, 937 F.2d 589, 19 USPQ2d 1289 (Fed. Cir. 1991). Consequently, the factual inquiries set forth in *Graham v. John Deere Co.,* 383 U.S. 1, 148 USPQ 459 (1966) are applied for establishing a background for determining obviousness are employed when making an obvious-type double patenting analysis. These factual inquiries are: (A) determine the scope and content of a patent claim relative to a claim in the application at issue; (B) determine the differences between the scope and content of the patent claim as determined in (A) and the claim in the application at issue; (C) determine the level of ordinary skill in the pertinent art; and (D) evaluate any objective indicia of nonobviousness. The conclusion of obviousness-type double patenting is made in light of these factual determinations, and any rejection should make clear: (A) the differences between the inventions defined by the conflicting claims; and (B) the reasons why a person of ordinary skill in the art would conclude that the invention defined in the claim at issue would have been an obvious variation of the invention defined in a claim in the patent. There is no indication in the record that the examiner has performed the required analysis. Furthermore, the examiner has not set forth the differences between the inventions defined by the conflicting

claims, and the reasons why a person of ordinary skill in the art would conclude that the invention defined in the claim at issue would have been an obvious variation of the invention defined in a claim in the patent. The rejections should therefore be reversed.

When the correct standard is applied the rejection fails because the pending claims in this application are patentably distinct from the claims of the '214 patent. Claim 1 of the '214 patent recites:

> A method of transmitting data in a data communications network, comprising the steps of: (i) establishing a connection-oriented communications link between a transmitter and a receiver through a Transmission Control Protocol (TCP) handshake, the communications link having a congestion window set to an initial length; (ii) transmitting data packets in TCP from the transmitter to the receiver; (iii) sending periodically a keep-alive request from the transmitter to the receiver, whereupon a re-transmission time-out timer is set, (iv) determining, at the transmitter, if an acknowledgment to the keep-alive request is not received before expiry of the re-transmission time-out timer, whereupon the transmitter backs off for a predetermined period; (v) detecting a missing data packet at the receiver; (vi) sending a negative acknowledgment from the receiver to the transmitter for the missing data packet, the receiver being unresponsive to any packets from the transmitter unless the receiver detects the missing data packet; (vii) decreasing, at the transmitter, the length of the congestion window in response to receipt of the negative acknowledgment; and (viii) re-transmitting the missing data packet.

Claim 1 of this application recites:

> A method for sending keep-alive messages by a node to a neighbor in a communication network, the method comprising: measuring a reliability of a communication link to the neighbor; and periodically calculating a reliability factor for communicating with a neighbor based upon the reliability of

> the communication link to the neighbor; and varying a frequency for sending keep-alive messages to the neighbor based upon the reliability factor.

In order to establish a *prima facie* case of obviousness the prior art references must teach or suggest all the claim limitations. (MPEP §2143). Claim 1 of the '214 patent fails to suggest measuring reliability of a communication link. Further, claim 1 of the '214 patent fails to suggest calculating a reliability factor for communicating with a neighbor based upon the reliability of the communication link, or varying the frequency for sending keep-alive messages to the neighbor based upon the reliability factor. Even if the cited prior art teaches or suggests the subject matter broadly recited in the pending claims, it does not suggest all of the limitations specifically recited in the claims. Appellant therefore requests that the rejections be reversed.

Conclusion

The rejections are improper for at least the reasons set forth above. Appellants accordingly request that the rejections be reversed and the application put forward for allowance.

Respectfully submitted,

/Holmes W. Anderson/
Holmes W. Anderson
Reg. No. 37,272
Attorney for Assignee

Date: September 6, 2011

Anderson Gorecki & Manaras LLP
33 Nagog Park
Acton MA 01720
(978) 264-4001

Appendix A - Claims

1. (previously presented) A method for sending keep-alive messages by a node to a neighbor in a communication network, the method comprising:

measuring a reliability of a communication link to the neighbor; and

periodically calculating a reliability factor for communicating with a neighbor based upon the reliability of the communication link to the neighbor;

varying a frequency for sending keep-alive messages to the neighbor based upon the reliability factor; and

sending keep-alive messages by the node to the neighbor in accordance with those steps.

2. (previously presented) The method of claim 1, wherein calculating the reliability factor for communicating with the neighbor comprises:

determining a reliability for the neighbor; and

wherein the step of calculating the reliability factor is further based upon the reliability for the neighbor.

3. (cancelled)

4. (previously presented) The method of claim 2, wherein calculating the reliability factor for communicating with the neighbor comprises:

determining a reliability for the neighbor;

measuring a reliability of a communication link to the neighbor;

assigning a relative weight to each of the reliability for the neighbor and the reliability of the communication link to the neighbor;

calculating the reliability factor to be a weighted average of the reliability for the neighbor and the reliability of the communication link to the neighbor.

5. (Previously Presented) The method of claim 1, wherein varying the frequency for sending keep-alive messages to the neighbor based upon the reliability comprises:

setting the frequency for sending keep-alive messages to the neighbor in inverse proportion to the reliability factor.

6. (Original) The method of claim 1, further comprising:

updating the reliability factor; and

adjusting the frequency for sending keep-alive messages to the neighbor based upon the reliability factor.

7. (Original) The method of claim 6, wherein adjusting the frequency for sending keep-alive messages to the neighbor comprises:

reducing the frequency for sending keep-alive messages to the neighbor, if the updated reliability factor represents a reliability improvement for communicating with the neighbor; and

increasing the frequency for sending keep-alive messages to the neighbor, if the updated reliability factor represents a reliability degradation for communicating with the neighbor.

8. (previously presented) A device for sending keep-alive message to a neighbor in a communication network, the device comprising:

a processor, a memory and a transmitter;

the processor executing a computer program stored in the memory, the computer program including:

reliability calculation logic operably coupled to determine a reliability for a communication link to the neighbor and periodically calculate a reliability factor for communicating with the neighbor based upon the reliability for the communication link to the neighbor; and

frequency variation logic responsive to the reliability calculation logic and operably coupled to calculate a frequency for sending keep-alive messages to the neighbor via the transmitter based upon the reliability factor.

9. (previously presented) The device of claim 8, wherein the reliability calculation logic is operably coupled to determine a reliability for the neighbor and wherein calculate the reliability factor is further calculated using based upon the reliability for the neighbor.

10. (cancelled)

11. (Previously Presented) The device of claim 8, wherein the reliability calculation logic is operably coupled to determine a reliability for the neighbor, measure a reliability for a communication link to the neighbor, assign a relative weight to each of the reliability for the neighbor and the reliability for the communication link to the neighbor, and calculate the reliability factor to be a weighted average of the reliability of the neighbor and the reliability of the communication link to the neighbor.

12. (Previously Presented) The device of claim 8, wherein the frequency variation logic is operably coupled to set the frequency for sending keep-alive messages to the neighbor in inverse proportion to the reliability factor.

13. (Previously Presented) The device of claim 8, wherein the reliability calculation logic is operably coupled to update the reliability factor, and wherein the frequency variation logic is operably coupled to adjust the frequency for sending keep-alive messages to the neighbor based upon the updated reliability factor.

14. (Previously Presented) The device of claim 13, wherein the frequency variation logic is operably coupled to reduce the frequency for sending keep alive messages to the neighbor if the updated reliability factor represents a reliability improvement for communicating with the neighbor and increase the frequency for sending keep-alive messages to the neighbor if the updated reliability factor represents a degradation for communicating with the neighbor.

15. (previously presented) A program product recorded on a computer readable medium for sending keep-alive messages to a neighbor in a communication network, the program product comprising:

reliability calculation logic operably coupled to measure a reliability for a communication link to the neighbor and to periodically calculate a reliability factor for communicating with the neighbor based upon the reliability for the communication link to the neighbor; and

frequency variation logic responsive to the reliability calculation logic and operably coupled to determine a frequency for sending keep-alive messages to the neighbor based upon the reliability factor.

16. (Previously Presented) The program product of claim 15, wherein the reliability calculation logic is programmed to determine a reliability for the neighbor and calculate the reliability factor based upon the reliability for the neighbor.

17. (cancelled)

18. (Previously Presented) The program product of claim 15, wherein the reliability calculation logic is programmed to determine a reliability for the neighbor, measure a reliability for a communication link to the neighbor, assign a relative weight to each of the reliability for the neighbor and the reliability for the communication link to the

neighbor, and calculate the reliability factor to be a weighted average of the reliability of the neighbor and the reliability of the communication link to the neighbor.

19. (Previously Presented) The program product of claim 15, wherein the frequency variation logic is programmed to set the frequency for sending keep-alive messages to the neighbor in inverse proportion to the reliability factor.

20. (Previously Presented) The program product of claim 15, wherein the reliability calculation logic is programmed to update the reliability factor, and wherein the frequency variation logic is operably coupled to adjust the frequency for sending keep-alive messages to the neighbor based upon the updated reliability factor.

21. (Previously Presented) The program product of claim 15, wherein the frequency variation logic is programmed to reduce the frequency for sending keep alive messages to the neighbor if the updated reliability factor represents a reliability improvement for communicating with the neighbor and increase the frequency for sending keep-alive messages to the neighbor if the updated reliability factor represents a degradation for communicating with the neighbor.

22. (previously presented) Apparatus comprising:

a plurality of interconnected devices including a node and a neighbor in communication over a link,

wherein the node is operably coupled to send keep-alive messages to the neighbor, and wherein the node is operably coupled to vary the frequency for sending keep-alive messages to the neighbor based upon a periodically computed reliability factor for communicating with the neighbor over the communication link,

wherein the node is operably coupled to calculate the reliability factor based upon a reliability for the neighbor and a measured reliability for the communication link.

23. (cancelled)

24. (previously presented) The apparatus of claim 22, wherein the node is operably coupled to set the frequency for sending keep-alive messages to the neighbor in inverse proportion to the reliability factor.

25. (previously presented) The method of claim 4, wherein the reliability factor (RF) is calculated using the below equation, where A is the measured reliability of the communication link to the neighbor, B is the determined reliability of the neighbor, W1 is a relative weight for A and W2 is a relative weight for B:

$$RF = (W1*A + W2*B).$$

26. (previously presented) The device of claim 11, wherein the reliability factor (RF) is calculated using the below equation, where A is the measured reliability of the communication link to the neighbor, B is the determined reliability of the neighbor, W1 is a relative weight for A and W2 is a relative weight for B:

$$RF = (W1*A + W2*B).$$

Appendix B - Evidence Submitted

None.

Appendix C - Related Proceedings

None.